arXiv:1404.5258v3 [math.CO] 12 Nov 2015

MAXIMUM-SIZE ANTICHAINS IN RANDOM SET-SYSTEMS

MAURÍCIO COLLARES NETO AND ROBERT MORRIS

Abstract. We show that, for $pn \to \infty$, the largest set in a p-random sub-family of the power set of $\{1, \dots, n\}$ containing no k-chain has size $(k-1+o(1))p\binom{n}{n/2}$ with high probability. This confirms a conjecture of Osthus.

1. Introduction

One of the cornerstones of extremal set theory is the famous theorem of Sperner [18], who proved in 1928 that the largest antichain in $\mathcal{P}(n)$, the family of all subsets of $\{1, \dots, n\}$, has size $\binom{n}{n/2}$. In 1945, Erdős [6] generalized this result by showing that any family of sets larger than the $k-1$ middle layers of $\mathcal{P}(n)$ contains a k-chain.

In this paper we will prove a sparse random analogue of Erdős' theorem. More precisely, for every function $p \gg 1/n$ we will determine, with high probability, the (asymptotic) size of the largest sub-family of $\mathcal{P}(n,p)$, the p-random sub-family[1] of $\mathcal{P}(n)$, containing no k-chain. This confirms a conjecture of Osthus [14].

Theorem 1.1. *Let $2 \leqslant k \in \mathbb{N}$, let $p = p(n)$ be such that $pn \to \infty$. Then the largest family $\mathcal{A} \subset \mathcal{P}(n,p)$ containing no k-chain has size*

$$|\mathcal{A}| = \big(k-1+o(1)\big)p\binom{n}{n/2} \tag{1}$$

with high probability as $n \to \infty$.

We remark that the bound on p is best possible, since the result fails to hold whenever $pn \to C$. Indeed, in this case Osthus [14] showed that, with high probability, the two middle layers of $\mathcal{P}(n,p)$ contain an antichain $\mathcal{A}$ of size $\big(1+e^{-C/2}+o(1)\big)p\binom{n}{n/2}$; adding $k-2$ further layers to $\mathcal{A}$ gives a family of size $\big(k-1+e^{-C/2}+o(1)\big)p\binom{n}{n/2}$ containing no k-chains.

The alert reader may have noticed that in order to obtain Theorem 1.1 it is sufficient to prove it in the case $k = 2$, since every family $\mathcal{A} \subset \mathcal{P}(n)$ containing no k-chain can be decomposed into $k-1$ antichains. (We thank Oliver Riordan for pointing this out to us.) Our proof will not use this simple fact, however, proceeding instead via an application of the hypergraph container method (see Section 2) to the k-uniform hypergraph encoding the k-chains of $\mathcal{P}(n)$. While more complicated that strictly necessary, this approach has two

Research supported in part by a CAPES bolsa Proex (MCN) and by CNPq Proc. 479032/2012-2 and Proc. 303275/2013-8 (RM).

[1] That is, $\mathcal{P}(n,p)$ is a random variable such that $\mathbb{P}(A \in \mathcal{P}(n,p)) = p$ for each $A \in \mathcal{P}(n)$, and such events are independent for different values of A.

significant advantages: it motivates the proof of Theorem 1.2, below, which we consider to be of independent interest, and it gives an easily accessible introduction to the approach[2] of [13], which we expect to have numerous other applications.

The study of the random set-system $\mathcal{P}(n,p)$ was initiated in 1961 by Rényi [15], who determined the threshold for the event that $\mathcal{P}(n,p)$ is an antichain. More recently, Kreuter [10] and Kohayakawa, Kreuter and Osthus [9] studied the length of the longest chain in $\mathcal{P}(n,p)$, and Kohayakawa and Kreuter [8] and Osthus [14] studied the size of the largest antichain. In particular, Osthus [14] proved that (1) holds if $pn \gg \log n$, and conjectured that $pn \gg 1$ is sufficient. We note that this conjecture has also been proved independently by Balogh, Mycroft and Treglown [2], who also studied sparser random set-systems.

The problem of obtaining sparse random analogues of classical results in extremal combinatorics has attracted a large amount of attention in recent years, culminating in the extraordinary breakthroughs of Conlon and Gowers [4] and Schacht [17], who developed general techniques for solving such problems, and as a result were able to prove a number of longstanding conjectures, such as sparse random analogues of the theorems of Turán and Szemerédi. A third approach, now known as the 'hypergraph container method', was subsequently developed independently by Balogh, Morris and Samotij [1] and by Saxton and Thomason [16]. We will use this latter method in order to prove Theorem 1.1.

In order to effectively apply the hypergraph container method (see Section 2), one requires a so-called 'balanced supersaturation theorem', and the proof of such a result (see Theorem 1.2, below) is the main innovation of this paper. An 'unbalanced' supersaturation theorem (giving a lower bound on the number of k-chains, but not controlling the distribution of these chains) was proved by Kleitman [7] in the case $k = 2$, and by Das, Gan and Sudakov [5] in general. More precisely, the authors of [5] used the permutation method pioneered by Katona and LYMB[3] in order to show that a family with t more elements than the extremal example above contains $\Omega\big(tn^{k-1}\big)$ k-chains. One of the key ideas from [5] will also play an important role in our proof, see Lemma 3.4 below.

In order to state our balanced supersaturation theorem, we will need a couple of simple definitions. For each $k \geqslant 2$ and $n \in \mathbb{N}$, let $\mathcal{G}_k = \mathcal{G}_k(n)$ denote the k-uniform hypergraph on vertex set $\mathcal{P}(n)$ whose edges encode k-chains, i.e., $\{F_1, \dots, F_k\} \in E(\mathcal{G}_k)$ if and only if $F_1 \supsetneq \cdots \supsetneq F_k$ for some ordering of the elements. Given $\mathcal{F} \subset \mathcal{P}(n)$, we write $\mathcal{H} \subset \mathcal{G}_k[\mathcal{F}]$ to denote that $\mathcal{H}$ is a k-uniform hypergraph with vertex set $\mathcal{F}$ whose edges are all members of $E(\mathcal{G}_k)$. For each $\ell \in [k]$, we write $\Delta_\ell(\mathcal{H})$ for the maximum degree of a set of ℓ vertices of $\mathcal{H}$, that is

$$\Delta_\ell(\mathcal{H}) = \max\big\{d_{\mathcal{H}}(\mathcal{A}) \,:\, \mathcal{A} \subset V(\mathcal{H}),\, |\mathcal{A}| = \ell\big\},$$

[2]This approach allows one to efficiently apply the hypergraph container lemma an unbounded number of times by combining it with a suitable 'balanced supersaturation theorem', such as Theorem 1.2. It was used in [13], to prove (amongst other things) that the number of C_{2k}-free graphs with n vertices is $2^{O(n^{1+1/k})}$.

[3]The acronym LYMB refers to Lubell [11], Yamamoto [19], Mešalkin [12] and Bollobás [3]. It often causes spelling confusion due to the silent B.

where $d_{\mathcal{H}}(\mathcal{A}) = \big|\big\{\, \mathcal{B} \in E(\mathcal{H}) : \mathcal{A} \subset \mathcal{B} \,\big\}\big|$. We also write $\mathcal{I}(\mathcal{H})$ for the collection of independent sets of $\mathcal{H}$, and $\alpha(\mathcal{H})$ for the size of the largest member of $\mathcal{I}(\mathcal{H})$.

We can now state the key new tool that we will use to prove Theorem 1.1. It says that a family with slightly more than $\alpha(\mathcal{G}_k) = \big(k-1+o(1)\big)\binom{n}{n/2}$ elements not only contains many k-chains, but that these chains can be chosen to be fairly 'evenly distributed' over $\mathcal{P}(n)$.

Theorem 1.2. *For every $k \geqslant 2$ and $\alpha > 0$, there exists $\delta = \delta(\alpha, k) > 0$ such that the following holds. Let $n \in \mathbb{N}$ and $\mathcal{F} \subset \mathcal{P}(n)$ satisfy $|\mathcal{F}| \geqslant (k-1+\alpha)\binom{n}{n/2}$, and suppose that $\delta^{-1} \leqslant m \leqslant \binom{|F|}{|G|}$ for every $F, G \in \mathcal{F}$ with $F \supsetneq G$. Then there exists $\mathcal{H} \subset \mathcal{G}_k[\mathcal{F}]$ satisfying*

(a) $e(\mathcal{H}) \geqslant \delta^k m^{k-1}\binom{n}{n/2}$,
(b) $\Delta_\ell(\mathcal{H}) \leqslant (\delta m)^{k-\ell}$ for every $1 \leqslant \ell \leqslant k$.

We remark that the bounds in Theorem 1.2 are all close to best possible. To see this, set $m = n/3$ and consider the $k-1$ middle layers of the hypercube, together with $\alpha\binom{n}{n/2}$ elements from the next layer up. Then $\mathcal{G}_k[\mathcal{F}]$ has $O\big(n^{k-1}\binom{n}{n/2}\big)$ edges and $\Delta_\ell\big(\mathcal{G}_k[\mathcal{F}]\big) = \Omega(n^{k-\ell})$ for every $1 \leqslant \ell \leqslant k$. The technical assumption $m \leqslant \binom{|F|}{|G|}$ for every $F, G \in \mathcal{F}$ with $F \supsetneq G$ will be useful because it will allow us to deduce sufficiently strong bounds both when $|\mathcal{F}|$ is close to $\alpha(\mathcal{G}_k)$, and when it is much larger, see Section 2.

The rest of this paper is organised as follows. In Section 2 we outline the hypergraph container method, in Section 3 we prove Theorem 1.2, and in Section 4 we perform the necessary technical computations in order to deduce Theorem 1.1.

2. Hypergraph containers

In this section, we will recall the powerful method of hypergraph containers, which was recently introduced in [1, 16]. Roughly speaking, the method allows us to deduce from a balanced supersaturation theorem (such as Theorem 1.2) that there exists a relatively small family of 'containers', each not too large, which cover the family $\mathcal{I}(\mathcal{H})$ of independent sets of a k-uniform hypergraph $\mathcal{H}$. The key container lemma (see [1, Proposition 3.1] and [16, Theorem 3.4]) is as follows.

The Hypergraph Container Lemma. *For every $k \in \mathbb{N}$ and $c > 0$, there exists a $\delta > 0$ such that the following holds. Let $\tau \in (0,1)$ and suppose that $\mathcal{H}$ is a k-uniform hypergraph on N vertices such that*

$$\Delta_\ell(\mathcal{H}) \leqslant c \cdot \tau^{\ell-1}\frac{e(\mathcal{H})}{N} \tag{2}$$

for every $1 \leqslant \ell \leqslant k$. Then there exist a family $\mathcal{C}$ of subsets of $V(\mathcal{H})$, and a function $f\colon \mathcal{P}\big(V(\mathcal{H})\big) \to \mathcal{C}$ such that:

(a) For every $I \in \mathcal{I}(\mathcal{H})$ there exists $T \subset I$ with $|T| \leqslant k \cdot \tau N$ and $I \subset f(T) \cup T$,
(b) $|C| \leqslant (1-\delta)N$ for every $C \in \mathcal{C}$.

We first apply this lemma to the hypergraph $\mathcal{G}_k$, to obtain a large family $\mathcal{C}_1$ of containers, each of size at most $(1-\delta)2^n$. We then apply the lemma again, for each $\mathcal{F} \in \mathcal{C}_1$ with $|\mathcal{F}| \geqslant (k-1+\alpha)\binom{n}{n/2}$ (for some small $\alpha > 0$), to the hypergraph $\mathcal{H} \subset \mathcal{G}_k[\mathcal{F}]$ given by Theorem 1.2. We repeat this process until all containers have size at most $(k-1+\alpha)\binom{n}{n/2}$. The conditions (a) and (b) in Theorem 1.2 allow us to check that (2) holds for a suitable value of τ, and hence to count the containers in our final collection. See [13] for a similar application of the container lemma in the context of C_{2k}-free graphs.

In order to further motivate the statement of Theorem 1.2 (and the technical condition $m \leqslant \binom{|F|}{|G|}$ for every $F, G \in \mathcal{F}$ with $F \supsetneq G$), we will next deduce from it the following two lemmas, which we will use to check the condition (2). The first shows that we can take $\tau = 1/n$ when $\mathcal{F}$ is slightly larger than $\alpha(\mathcal{G}_k)$.

Lemma 2.1. *For every $k \geqslant 2$ and $\alpha > 0$, there exists $c = c(\alpha, k) > 0$ such that the following holds. Let $n \in \mathbb{N}$ be sufficiently large and $\mathcal{F} \subset \mathcal{P}(n)$ satisfy $(k-1+\alpha)\binom{n}{n/2} \leqslant |\mathcal{F}| \leqslant 3k\binom{n}{n/2}$. Then there exists $\mathcal{H} \subset \mathcal{G}_k[\mathcal{F}]$ satisfying*

$$\Delta_\ell(\mathcal{H}) \leqslant \frac{c}{n^{\ell-1}} \cdot \frac{e(\mathcal{H})}{|\mathcal{F}|}$$

for every $1 \leqslant \ell \leqslant k$.

Proof. First, observe that (by adjusting α slightly) we may assume that $|F| \geqslant n/3$ for every $F \in \mathcal{F}$, since the number of sets smaller than this is much smaller than $\binom{n}{n/2}$. Thus, applying Theorem 1.2 with $m = n/3$, it follows that there exists a hypergraph $\mathcal{H} \subset \mathcal{G}_k[\mathcal{F}]$ and a constant $\delta = \delta(\alpha, k) > 0$ with $e(\mathcal{H}) \geqslant \delta^k m^{k-1}\binom{n}{n/2}$ and $\Delta_\ell(\mathcal{H}) \leqslant (\delta m)^{k-\ell}$ for every $1 \leqslant \ell \leqslant k$. It follows that

$$\Delta_\ell(\mathcal{H}) \leqslant (\delta m)^{k-\ell} = \frac{3^\ell k}{\delta^\ell n^{\ell-1}} \cdot \frac{\delta^k m^{k-1}\binom{n}{n/2}}{3k\binom{n}{n/2}} \leqslant \frac{c}{n^{\ell-1}} \cdot \frac{e(\mathcal{H})}{|\mathcal{F}|},$$

where $c = k \cdot (3/\delta)^k$, as required. □

The next lemma shows that if $|\mathcal{F}|$ is larger, then we can in fact take τ much smaller.

Lemma 2.2. *For every $k \geqslant 2$, there exists $c = c(k) > 0$ such that the following holds. Let $n \in \mathbb{N}$ be sufficiently large and $\mathcal{F} \subset \mathcal{P}(n)$ satisfy $|\mathcal{F}| \geqslant 3k\binom{n}{n/2}$. Then there exists $\mathcal{H} \subset \mathcal{G}_k[\mathcal{F}]$ satisfying*

$$\Delta_\ell(\mathcal{H}) \leqslant \frac{c}{n^{3\ell-3}} \cdot \frac{e(\mathcal{H})}{|\mathcal{F}|}$$

for every $1 \leqslant \ell \leqslant k$.

Proof. First, choose an arbitrary partition $\mathcal{F} = \mathcal{F}_0 \cup \mathcal{F}_1 \cup \cdots \cup \mathcal{F}_t$ such that $|\mathcal{F}_i| = 3k\binom{n}{n/2}$ for every $i \in [t]$ and $|\mathcal{F}_0| < 3k\binom{n}{n/2}$. Fix $i \in [t]$, and observe that, by the pigeonhole principle, there are at least $k\binom{n}{n/2}$ elements of $\mathcal{F}_i$ whose sizes have the same remainder modulo 3. Let $\mathcal{F}_i'$ be a collection of $\big(k - o(1)\big)\binom{n}{n/2}$ such elements, all of size at least $n/3$, and note that $\binom{|F|}{|G|} \geqslant \binom{n/3}{3}$ for every $F, G \in \mathcal{F}_i'$ with $F \supsetneq G$. Thus, applying Theorem 1.2 with $m = \binom{n/3}{3}$,

it follows that there exists a hypergraph $\mathcal{H}_i \subset \mathcal{G}_k[\mathcal{F}_i']$ and a constant $\delta = \delta(k) > 0$ such that $e(\mathcal{H}_i) = \lceil \delta^k m^{k-1} \binom{n}{n/2} \rceil$ and

$$\Delta_\ell(\mathcal{H}_i) \leqslant (\delta m)^{k-\ell} \leqslant \frac{c'}{n^{3\ell-3}} \cdot \frac{\delta^k m^{k-1} \binom{n}{n/2}}{k\binom{n}{n/2}} \leqslant \frac{c'}{n^{3\ell-3}} \cdot \frac{e(\mathcal{H}_i)}{|\mathcal{F}_i'|}$$

for some $c' = c'(k)$ and every $1 \leqslant \ell \leqslant k$. Setting $\mathcal{H} = \mathcal{H}_1 \cup \cdots \cup \mathcal{H}_t$, it follows that

$$\Delta_\ell(\mathcal{H}) \leqslant \max_{1 \leqslant i \leqslant t} \big\{\Delta_\ell(\mathcal{H}_i)\big\} \leqslant \frac{c'}{n^{3\ell-3}} \cdot \frac{\max_i e(\mathcal{H}_i)}{\min_i |\mathcal{F}_i'|} \leqslant \frac{c}{n^{3\ell-3}} \cdot \frac{e(\mathcal{H})}{|\mathcal{F}|}$$

as claimed, since $e(\mathcal{H}) = \sum_{i=1}^t e(\mathcal{H}_i) = t \cdot e(\mathcal{H}_i)$ and $|\mathcal{F}| \leqslant 7t \cdot |\mathcal{F}_i'|$ for every $i \in [t]$. □

Motivated by the above bounds, fix $\tau \colon \mathcal{P}(n) \to \mathbb{R}$ to be the function defined by

$$\tau(A) := \begin{cases} n^{-1} & \text{if } |A| \leqslant 3k\binom{n}{n/2} \\ n^{-3} & \text{otherwise.} \end{cases} \tag{3}$$

We can now specialize the Hypergraph Container Lemma to our application by combining it with Lemma 2.1 and Lemma 2.2. The following corollary will be used in Section 4 to count the containers of a given size produced by repeated applications of the Hypergraph Container Lemma, see Theorem 4.2.

Corollary 2.3. *For every $2 \leqslant k \in \mathbb{N}$ and $\alpha > 0$, there exists $\delta = \delta(\alpha, k) > 0$ such that the following holds. Let $n \in \mathbb{N}$ be sufficiently large and $C \subset \mathcal{P}(n)$ with $|C| \geqslant (k - 1 + \alpha)\binom{n}{n/2}$. Then there exists a collection $\mathcal{C} \subset \mathcal{P}(C)$ and a function $f \colon \mathcal{P}(C) \to \mathcal{C}$ such that*

(a) For every $I \in \mathcal{I}(\mathcal{G}_k[C])$, there exists T with $|T| \leqslant k \cdot \tau(C)|C|$ and $T \subset I \subset f(T) \cup T$.

(b) $|C'| \leqslant (1 - \delta)|C|$ for every $C' \in \mathcal{C}$.

Proof. Apply the Hypergraph Container Lemma to the hypergraph $\mathcal{H} \subset \mathcal{G}_k[C]$ given by Lemma 2.1 (if $|C| \leqslant 3k\binom{n}{n/2}$), or by Lemma 2.2 (otherwise), and observe that (for a suitable choice of the constant c) the inequality (2) holds with $\tau = \tau(C)$ for every $1 \leqslant \ell \leqslant k$. It follows immediately that there exist a family $\mathcal{C}$ of subsets of C, and a function $f \colon \mathcal{P}(C) \to \mathcal{C}$ such that (a) and (b) hold, as required. □

3. Balanced supersaturation

In this section, we will prove Theorem 1.2 by constructing $\mathcal{H}$ one edge at a time. More precisely, starting with $\mathcal{H} = \emptyset$, we will repeatedly apply the following lemma, adding new edges to $\mathcal{H}$ until the conditions of Theorem 1.2 are satisfied.

Lemma 3.1. *For every $k \geqslant 2$ and $\alpha > 0$, there exists $\delta = \delta(\alpha, k) > 0$ such that the following holds. Let $n \in \mathbb{N}$ and $\mathcal{F} \subset \mathcal{P}(n)$ satisfy $|\mathcal{F}| \geqslant (k - 1 + \alpha)\binom{n}{n/2}$, and suppose that $\delta^{-1} \leqslant m \leqslant \binom{|F|}{|G|}$ for every $F, G \in \mathcal{F}$ with $F \supsetneq G$. If $\mathcal{H} \subset \mathcal{G}_k[\mathcal{F}]$ is a hypergraph satisfying*

(a) $e(\mathcal{H}) \leqslant \delta^k m^{k-1}\binom{n}{n/2}$,

(b) $\Delta_\ell(\mathcal{H}) \leqslant (\delta m)^{k-\ell}$ for every $\ell \in [k]$,

then there exists an edge $f \in \mathcal{G}_k[\mathcal{F}] \setminus \mathcal{H}$ for which $\Delta_\ell(\{f\} \cup \mathcal{H}) \leqslant (\delta m)^{k-\ell}$ for every $\ell \in [k]$.

We remark that this approach to proving balanced supersaturation theorems – adding the edges of $\mathcal{H}$ one by one – was also used in [13], and is likely to have further applications.

The rest of this section will be dedicated to proving the Lemma 3.1, so from now on let us fix $\alpha > 0$ and $k \geqslant 2$, and choose $\delta > 0$ sufficiently small and $m \geqslant \delta^{-1}$. Moreover, let us fix $n \in \mathbb{N}$, a family $\mathcal{F} \subset \mathcal{P}(n)$ and a hypergraph $\mathcal{H} \subset \mathcal{G}_k[\mathcal{F}]$ satisfying the conditions of the lemma. The degree function of $\mathcal{H}$ will simply be denoted by d, for simplicity.

We say that a non-empty family $\mathcal{A} \subset \mathcal{P}(n)$ is *saturated* if $d(\mathcal{A}) = \lfloor (\delta m)^{k-|\mathcal{A}|} \rfloor$, that is, if no edge of $\mathcal{F}$ containing this family can be added to the hypergraph $\mathcal{H}$ without violating condition (b). A family $\mathcal{B} \subset \mathcal{P}(n)$ is *bad* if it contains a saturated sub-family $\mathcal{A} \subset \mathcal{B}$. Otherwise we say that $\mathcal{B}$ is *good*. With this terminology, the conclusion of Lemma 3.1 is that $\mathcal{G}_k[\mathcal{F}]$ contains a good edge. Indeed, since a good edge $f \in \mathcal{G}_k[\mathcal{F}]$ is not saturated, then $d(f) < 1$, and so $f \notin \mathcal{H}$.

The following easy lemma will be a crucial tool in the proof of Lemma 3.1. It says that there are not too many ways to turn a good family bad.

Lemma 3.2. *For any good $\mathcal{A} \subset \mathcal{P}(n)$, there are at most $2^{|\mathcal{A}|} \cdot 2\delta k m$ sets $F \in \mathcal{P}(n)$ for which $\{F\} \cup \mathcal{A}$ is bad and $\{F\}$ is not saturated.*

Proof. The result follows from a simple double-counting argument, which we spell out below. Since $\mathcal{A}$ is good, any saturated sub-family of $\{F\} \cup \mathcal{A}$ must contain F. In other words, any F such that $\{F\} \cup \mathcal{A}$ is bad belongs to

$$S(\mathcal{B}) := \big\{ F \in \mathcal{P}(n) \,:\, d\big(\{F\} \cup \mathcal{B}\big) = \big\lfloor (\delta m)^{k-|\mathcal{B}|-1} \big\rfloor \big\}$$

for some $\mathcal{B} \subset \mathcal{A}$. Moreover, if $\{F\}$ is not saturated, then $\mathcal{B}$ cannot be empty. Therefore, it is enough to bound the size of $S(\mathcal{B})$ when $\mathcal{B}$ is non-empty. We do so by noting that

$$|S(\mathcal{B})| \big\lfloor (\delta m)^{k-|\mathcal{B}|-1} \big\rfloor = \sum_{F \in S(\mathcal{B})} d\big(\{F\} \cup \mathcal{B}\big) \leqslant k d(\mathcal{B}) \leqslant k(\delta m)^{k-|\mathcal{B}|},$$

where the first inequality is true because each edge of $\mathcal{H}$ containing $\mathcal{B}$ contributes at most k to the sum. Since $m \geqslant \delta^{-1}$, we obtain $|S(\mathcal{B})| \leqslant 2\delta k m$. The claimed bound now follows by summing over all choices of $\mathcal{B}$. □

Similarly, noting that $S(\emptyset) = \big\{ F \in \mathcal{P}(n) : \{F\} \text{ is saturated} \big\}$, we have

$$|S(\emptyset)| \lfloor (\delta m)^{k-1} \rfloor \leqslant k \cdot e(\mathcal{H}).$$

By condition (a) and the bound $m \geqslant \delta^{-1}$, it follows that $|S(\emptyset)| \leqslant 2\delta k \binom{n}{n/2}$. Thus, by adjusting α slightly if necessary, we can remove the elements of S from $\mathcal{F}$. Therefore, from now on we will assume that $\mathcal{F}$ contains no saturated sets.

We will next sketch the proof of Lemma 3.1. The key idea is that if we choose F_1 to be of *minimal cardinality* such that the "density" of k-chains below F_1 (see Definition 3.3) is bigger than α/k (see Lemma 3.6), then only few of those k-chains will be bad, and hence at least one of them will be good. In order to bound the density of bad k-chains below F_1, let us define

a chain $F_1 \supsetneq \cdots \supsetneq F_\ell$ to be *critical* if $\{F_1, \dots, F_{\ell-1}\}$ is good but $\{F_1, \dots, F_\ell\}$ is not. We will use Lemma 3.2 to show that the density of critical ℓ-chains is small (see Lemma 3.7). We will then use the minimality of F_1 to deduce that the operation of extending critical ℓ-chains to bad k-chains only increases the density by a bounded factor.

In order to make the above sketch more precise, let us next formalize the notion of density that we will use. This definition is inspired by the work of Das, Gan and Sudakov [5], see Lemma 3.4 below. We remark that, despite its name, the ℓ-chain density of a set is not bounded above by 1, and in fact can be as large as $\Omega(n^{\ell-1})$.

Definition 3.3. The ℓ-chain density of a set $F_1 \in \mathcal{F}$, denoted by $\mathrm{c}_\ell(F_1)$, is given by

$$\mathrm{c}_\ell(F_1) := \sum_{\substack{F_2,\dots,F_\ell \in \mathcal{F} \\ F_1 \supsetneq F_2 \supsetneq \cdots \supsetneq F_\ell}} \binom{|F_1|}{|F_2|}^{-1} \cdots \binom{|F_{\ell-1}|}{|F_\ell|}^{-1}$$

In particular, $\mathrm{c}_1(F) = 1$ for all $F \in \mathcal{F}$.

The following lemma is essentially due to Das, Gan and Sudakov [5]. Since it was not explicitly stated in their paper, we will give the proof for completeness.

Lemma 3.4 (Das, Gan and Sudakov)**.** *For any fixed $1 \leqslant i < j \leqslant k$, we have*

$$\sum_{F \in \mathcal{F}} \frac{1}{\binom{n}{|F|}} \big(\mathrm{c}_i(F) - \mathrm{c}_j(F) \big) \leqslant \max_{s \in \mathbb{N}} \binom{s}{i} - \binom{s}{j}.$$

Proof. Following the permutation method, say a permutation π of $[n]$ *contains* a set F if $F = \{\pi(1), \dots, \pi(|F|)\}$. Moreover, say it contains a chain if it contains all sets of the chain. Note that the number of permutations containing a given chain $F_1 \supsetneq \cdots \supsetneq F_\ell$ is

$$(n - |F_1|)! \cdot |F_1 \setminus F_2|! \cdots |F_{\ell-1} \setminus F_\ell|! \cdot |F_\ell|! = n! \cdot \binom{n}{|F_1|}^{-1} \binom{|F_1|}{|F_2|}^{-1} \cdots \binom{|F_{\ell-1}|}{|F_\ell|}^{-1},$$

and so, denoting by $X_\ell(\pi)$ the number of ℓ-chains contained in π, the expected value of X_ℓ with respect to the uniform probability measure on the set of permutations is

$$\mathbb{E}(X_\ell) = \sum_{\substack{F_1,\dots,F_\ell \in \mathcal{F} \\ F_1 \supsetneq \dots \supsetneq F_\ell}} \binom{n}{|F_1|}^{-1} \binom{|F_1|}{|F_2|}^{-1} \cdots \binom{|F_{\ell-1}|}{|F_\ell|}^{-1} = \sum_{F_1 \in \mathcal{F}} \mathrm{c}_\ell(F_1) \Big/ \binom{n}{|F_1|}.$$

On the other hand, since sets contained in a single permutation always form a chain, $X_\ell(\pi)$ equals $\binom{s}{\ell}$, where s is the number of elements of $\mathcal{F}$ contained in π. We deduce that

$$X_i(\pi) - X_j(\pi) \leqslant \max_{s \in \mathbb{N}} \binom{s}{i} - \binom{s}{j},$$

and the conclusion follows by taking the expected value of both sides. □

A very useful feature of Lemma 3.4 is that the upper bound it provides does not depend on n. We will next use this to show that ℓ-chain densities cannot decrease too quickly as a function of ℓ, and hence that it is enough to upper bound the k-chain density of a set whenever we want an upper bound for all of its lower densities.

Lemma 3.5. *For every $F \in \mathcal{F}$ and $1 \leqslant \ell < k$, we have $\mathrm{c}_\ell(F) \leqslant \mathrm{c}_k(F) + 4^k$.*

Proof. The result is trivial for $\ell = 1$, as $\mathrm{c}_1(F) = 1$. For $\ell \geqslant 2$, we can use the identity

$$\mathrm{c}_\ell(F) = \sum_{\substack{F_2 \in \mathcal{F} \\ F \supsetneq F_2}} \binom{|F|}{|F_2|}^{-1} \sum_{\substack{F_3,\dots,F_k \in \mathcal{F} \\ F_2 \supsetneq \cdots \supsetneq F_\ell}} \binom{|F_2|}{|F_3|}^{-1} \cdots \binom{|F_{\ell-1}|}{|F_\ell|}^{-1} = \sum_{F \supsetneq F_2 \in \mathcal{F}} \mathrm{c}_{\ell-1}(F_2) \Big/ \binom{|F|}{|F_2|}$$

together with Lemma 3.4 (applied to the hypercube of subsets of F) to obtain

$$\mathrm{c}_\ell(F) - \mathrm{c}_k(F) = \sum_{F \supsetneq F_2 \in \mathcal{F}} \frac{1}{\binom{|F|}{|F_2|}} (\mathrm{c}_{\ell-1}(F_2) - \mathrm{c}_{k-1}(F_2)) \leqslant \max_{s \in \mathbb{N}} \binom{s}{\ell-1} - \binom{s}{k-1}.$$

Since the function being maximized is negative for all $s \geqslant 2k-1$, the right side is at most $\binom{2k-1}{\ell-1} \leqslant 4^k$, which proves the result. □

Lemma 3.4 also allows us to deduce that at least one element of our family has large k-chain density, as we show in the following pigeonhole-like observation.

Lemma 3.6. *If $0 \leqslant \alpha \leqslant 1$ and $|\mathcal{F}| \geqslant (k-1+\alpha)\binom{n}{n/2}$, then $\max_F \mathrm{c}_k(F) \geqslant \alpha/k$.*

Proof. By Lemma 3.4 with $i = 1$ and $j = k$, and since $c_1(F) = 1$, we have

$$\sum_{F \in \mathcal{F}} \frac{1}{\binom{n}{|F|}} (1 - \mathrm{c}_k(F)) \leqslant \max_{s \in \mathbb{N}} \binom{s}{1} - \binom{s}{k} = k - 1.$$

However, if the desired conclusion were not true, we would have

$$\sum_{F \in \mathcal{F}} \frac{1}{\binom{n}{|F|}} (1 - \mathrm{c}_k(F)) > \sum_{F \in \mathcal{F}} \frac{1}{\binom{n}{n/2}} \left(1 - \frac{\alpha}{k}\right) \geqslant (k - 1 + \alpha) \cdot \frac{k - \alpha}{k} \geqslant k - 1,$$

where, for the last step, note that equality holds when $\alpha \in \{0, 1\}$. □

Finally, we will need the following lemma, which bounds the density of critical ℓ-chains. It is a simple consequence of Lemma 3.2 and our assumption that $m \leqslant \binom{|F|}{|G|}$ for every $F, G \in \mathcal{F}$ with $F \supsetneq G$.

Lemma 3.7. *For every $F_1 \in \mathcal{F}$ and $1 \leqslant \ell < k$,*

$$\sum_{\substack{F_2,\dots,F_{\ell+1} \in \mathcal{F} \\ F_1 \supsetneq \cdots \supsetneq F_{\ell+1} \text{ critical}}} \binom{|F_1|}{|F_2|}^{-1} \cdots \binom{|F_\ell|}{|F_{\ell+1}|}^{-1} \leqslant 2^\ell \cdot 2\delta k \cdot \mathrm{c}_\ell(F_1) \tag{4}$$

Proof. Recall that if $F_1 \supsetneq \cdots \supsetneq F_{\ell+1}$ is critical, then $\{F_1, \dots, F_\ell\}$ is good but $\{F_1, \dots, F_{\ell+1}\}$ is not. By Lemma 3.2, it follows that the left-hand side of (4) is at most

$$\sum_{\substack{F_2,\dots,F_\ell \in \mathcal{F} \\ F_1 \supsetneq \cdots \supsetneq F_\ell}} \binom{|F_1|}{|F_2|}^{-1} \cdots \binom{|F_{\ell-1}|}{|F_\ell|}^{-1} \cdot 2^\ell \cdot 2\delta k m \cdot \max_{F_\ell \supsetneq F_{\ell+1} \in \mathcal{F}} \binom{|F_\ell|}{|F_{\ell+1}|}^{-1}.$$

The result then follows from our upper bound on m and the definition of $c_\ell(F_1)$. □

We are now ready to carry out the plan outlined above, and prove Lemma 3.1.

Proof of Lemma 3.1. We may assume, without loss of generality, that $0 < \alpha < 1$. Let F_1 be of minimal cardinality such that $\mathrm{c}_k(F_1) \geqslant \alpha/k$ (note that at least one such F_1 exists, by Lemma 3.6). We claim that

$$\sum_{\substack{F_2,\dots,F_k \in \mathcal{F} \\ F_1 \supsetneq \cdots \supsetneq F_k \text{ bad}}} \binom{|F_1|}{|F_2|}^{-1} \cdots \binom{|F_{k-1}|}{|F_k|}^{-1} \leqslant \frac{\mathrm{c}_k(F_1)}{2}, \tag{5}$$

which immediately implies that the total k-chain density of *good* chains is positive, and therefore that at least one good chain exists. In order to prove (5), notice that every bad k-chain $F_1 \supsetneq \cdots \supsetneq F_k$ is associated with a unique $1 \leqslant \ell < k$ such that $F_1 \supsetneq \cdots \supsetneq F_{\ell+1}$ is critical. As such, we can write the left side of (5) as

$$\sum_{\ell=1}^{k-1} \left(\sum_{\substack{F_2,\dots,F_{\ell+1} \in \mathcal{F} \\ F_1 \supsetneq \cdots \supsetneq F_{\ell+1} \text{ critical}}} \binom{|F_1|}{|F_2|}^{-1} \cdots \binom{|F_\ell|}{|F_{\ell+1}|}^{-1} \cdot \mathrm{c}_{k-\ell}(F_{\ell+1}) \right).$$

We will proceed by bounding each term of the outer sum separately, so fix $1 \leqslant \ell < k$. By Lemma 3.5 and the minimality of F_1, we have $\mathrm{c}_{k-\ell}(F_{\ell+1}) \leqslant \mathrm{c}_k(F_{\ell+1}) + 4^k < \alpha/k + 4^k < 5^k$. Using this bound and Lemma 3.7, we obtain

$$\sum_{\substack{F_2,\dots,F_{\ell+1} \in \mathcal{F} \\ F_1 \supsetneq \cdots \supsetneq F_{\ell+1} \text{ critical}}} \binom{|F_1|}{|F_2|}^{-1} \cdots \binom{|F_\ell|}{|F_{\ell+1}|}^{-1} \cdot \mathrm{c}_{k-\ell}(F_{\ell+1}) \leqslant 2^\ell \cdot 2\delta k \cdot \mathrm{c}_\ell(F_1) \cdot 5^k. \tag{6}$$

Using Lemma 3.5 once again for the bound $\mathrm{c}_\ell(F_1) \leqslant \mathrm{c}_k(F_1) + 4^k$ and summing (6) over $1 \leqslant \ell < k$, we conclude that

$$\sum_{\substack{F_2,\dots,F_k \in \mathcal{F} \\ F_1 \supsetneq \cdots \supsetneq F_k \text{ bad}}} \binom{|F_1|}{|F_2|}^{-1} \cdots \binom{|F_{k-1}|}{|F_k|}^{-1} = \delta \cdot 2^{O(k)} \cdot (\mathrm{c}_k(F_1) + 4^k) = \frac{\delta \cdot 2^{O(k)}}{\alpha} \cdot \mathrm{c}_k(F_1),$$

since $\mathrm{c}_k(F_1) \geqslant \alpha/k$. The right side can be made less than $\mathrm{c}_k(F_1)/2$ by choosing δ to be small (only as a function of α and k), and so the proof is complete. □

4. Proof of Theorem 1.1

In this section we will deduce Theorem 1.1 from the results of the previous two sections. More precisely, we will use Corollary 2.3 to prove a 'fingerprint theorem' (Theorem 4.2, below), which easily implies Theorem 1.1. A *coloured vertex set* is simply a family $\mathcal{A} \subset \mathcal{P}(n)$ together with a function $c\colon \mathcal{A} \to \mathbb{N}$. Recall that $\mathcal{G}_k$ denotes the k-uniform hypergraph whose edges encode k-chains. We will need the following definition.

Definition 4.1. A fingerprint of $\mathcal{G}_k$ is a family $\mathcal{S}$ of coloured vertex sets, together with:

(*a*) A fingerprint function $T\colon \mathcal{I}(\mathcal{G}_k) \to \mathcal{S}$ with $T(I) \subset I$ for every $I \in \mathcal{I}(\mathcal{G}_k)$.
(*b*) A container function $C\colon \mathcal{S} \to \mathcal{P}(V(\mathcal{G}_k))$ such that $I \subset C(T(I))$ for every $I \in \mathcal{I}(\mathcal{G}_k)$.

Each $S \in \mathcal{S}$ should be thought of as a sequence of subsets of $V(\mathcal{H})$ given by repeated application of the Hypergraph Container Lemma. The container function is obtained by applying the sequence of functions f given by these repeated applications. We will prove the following theorem.

Theorem 4.2. *For every $k \geqslant 2$ and $\varepsilon > 0$, there exist a constant $K = K(\varepsilon, k) > 0$ and a fingerprint $(\mathcal{S}, T, C)$ of $\mathcal{G}_k$ such that the following hold:*

(a) *Every $S \in \mathcal{S}$ satisfies $|S| \leqslant \frac{K}{n}\binom{n}{n/2}$;*

(b) *The number of members of $\mathcal{S}$ of size s is at most*

$$\left(\frac{K\binom{n}{n/2}}{s}\right)^s \exp\left(\frac{K}{n}\binom{n}{n/2}\right);$$

(c) $|C(T(I))| \leqslant (k-1+\varepsilon)\binom{n}{n/2}$ *for every $I \in \mathcal{I}(\mathcal{G}_k)$.*

Before proving Theorem 4.2, let us see how it implies Theorem 1.1.

Proof of Theorem 1.1. Let $k \geqslant 2$ and $\varepsilon > 0$ be arbitrary, and let $K = K(\varepsilon, k) > 0$ and $(\mathcal{S}, T, C)$ be the constant and fingerprint given by Theorem 4.2. Let $n \in \mathbb{N}$ be sufficiently large, and note that $pn \geqslant K\varepsilon^{-1}$, since $pn \to \infty$. If $I \subset \mathcal{P}(n,p)$ is an independent set of $\mathcal{G}_k$ of size at least $(k-1+3\varepsilon)p\binom{n}{n/2}$, then it follows that $T(I) \subset \mathcal{P}(n,p)$ and

$$\big|C(T(I)) \cap \mathcal{P}(n,p)\big| \geqslant \big(k-1+3\varepsilon\big)p\binom{n}{n/2}.$$

Let X be the number of elements of $\mathcal{S}$ for which these two properties hold. Then

$$\mathbb{E}(X) \leqslant \sum_{A \in \mathcal{S}} \mathbb{P}\big(A \subset \mathcal{P}(n,p)\big) \cdot \mathbb{P}\left(\big|(C(A) \setminus A) \cap \mathcal{P}(n,p)\big| \geqslant (k-1+2\varepsilon)p\binom{n}{n/2}\right),$$

where we used that $|A| \leqslant \varepsilon p\binom{n}{n/2}$ by the lower bound on pn and Theorem 4.2 (a). Hence, by the properties of $(\mathcal{S}, T, C)$ guaranteed by Theorem 4.2, and Chernoff's inequality,

$$\begin{aligned}\mathbb{E}(X) &\leqslant \sum_{s=1}^{\frac{K}{n}\binom{n}{n/2}} \left(\frac{K\binom{n}{n/2}}{s}\right)^s \exp\left(\frac{K}{n}\binom{n}{n/2}\right) \cdot p^s \cdot \exp\left(-\frac{\varepsilon^2 p}{3}\binom{n}{n/2}\right) \\ &\leqslant \frac{K}{n}\binom{n}{n/2} \exp\left(\frac{K\log(pn)}{n}\binom{n}{n/2} + \frac{K}{n}\binom{n}{n/2} - \frac{\varepsilon^2 p}{3}\binom{n}{n/2}\right),\end{aligned}$$

since the summand is increasing in s on the interval $\big(0, (Kp/e)\binom{n}{n/2}\big)$, and $K/n \ll Kp/e$. Therefore, by Markov's inequality, and since $pn \gg \log(pn) \gg 1$, we have

$$\mathbb{P}\left(\alpha\big(\mathcal{P}(n,p)\big) \geqslant \big(k-1+3\varepsilon\big)p\binom{n}{n/2}\right) \leqslant \exp\left(-\frac{\varepsilon^2 p}{6}\binom{n}{n/2}\right) \to 0$$

as $n \to \infty$, as required. □

It only remains to prove Theorem 4.2. We will use a straightforward but technical lemma.

Lemma 4.3. *Let $M > 0$, $s > 0$ and $0 < \delta < 1$. For any finite sequence $(a_1, \dots, a_m)$ of real numbers summing to s such that $1 \leqslant a_j \leqslant (1-\delta)^j M$ for each $j \in [m]$, we have*

$$s \log s \leqslant \sum_{j=1}^{m} a_j \log a_j + O(M).$$

We remark that this lemma is especially easy to prove if $m = O(1)$, which will be the case in our application. However, it is not much harder to prove in general, and this more general version is necessary for other applications, cf. [13, Section 6].

Proof of Lemma 4.3. Fix $m \in \mathbb{N}$ and note that, by compactness, we can assume that the sequence $(a_1, \dots, a_m)$ achieves the minimum of $\sum_{j=1}^m x_j \log x_j$ subject to the given conditions. Let

$$J_1 = \{j \in [m] : a_j < (1-\delta)^j M\}$$

and $J_2 = [m] \setminus J_1$; define also $s_i = \sum_{j \in J_i} a_j$ for $i \in \{1,2\}$. The convexity of $x \log x$ implies that all of the elements of the subsequence $(a_j)_{j \in J_1}$ are equal and that $J_1 = [t]$ for some $t \in \{0, \dots, m\}$, so that $s_1 \leqslant t(1-\delta)^t M$. Note that $s = \sum_j a_j = O(M)$ and

$$s_2 \log M - \sum_{j \in J_2} a_j \log a_j = \sum_{j \in J_2} a_j \log \frac{M}{a_j} \leqslant \sum_{j=1}^{\infty} (1-\delta)^j M \log \frac{1}{(1-\delta)^j} = O(M).$$

We are done if $t = 0$, so assume $t \geqslant 1$. By convexity, $s \log s \leqslant s_1 \log s_1 + s_2 \log s_2 + s \log 2$. Hence, recalling that $a_1 = \ldots = a_t = s_1/t$, we have

$$\begin{aligned} s \log s &\leqslant s_1 \log \frac{s_1}{t} + s_1 \log t + s_2 \log s_2 + O(M) \\ &\leqslant \sum_{j \in J_1} a_j \log a_j + t(1-\delta)^t M \log t + \sum_{j \in J_2} a_j \log a_j + O(M) \\ &= \sum_{j=1}^{m} a_j \log a_j + O(M), \end{aligned}$$

as claimed. □

We are now ready to prove the 'fingerprint theorem', and thus complete the proof of Theorem 1.1.

Proof of Theorem 4.2. Let $k \geqslant 2$ and $\varepsilon > 0$ be arbitrary, let $\delta = \delta(\varepsilon, k) > 0$ be given by Corollary 2.3, choose a large constant $K = K(\varepsilon, k, \delta)$, and let $n \in \mathbb{N}$ be sufficiently large. For a given $I \in \mathcal{I}(\mathcal{G}_k)$, we will apply Corollary 2.3 a certain number of times, which we will denote by $m = m(I)$, to construct two sequences $C_1, \dots, C_{m+1}$ and $T_1, \dots, T_m$ of subsets of $V(\mathcal{G}_k)$. The construction will inductively maintain the following properties:

- (*i*) $I \subset C_{i+1} \cup T_1 \cup \cdots \cup T_i$,
- (*ii*) The sets $C_{i+1}, T_1, \dots, T_i$ are pairwise disjoint,
- (*iii*) C_{i+1} only depends on C_i and T_i,
- (*iv*) $|C_{i+1}| \leqslant (1-\delta)|C_i|$.

To do this, first set $C_1 := \mathcal{P}(n)$. As long as $|C_i| \geqslant (k-1+\varepsilon)\binom{n}{n/2}$, let $T_i \subset I \cap C_i$ and f_i be given by Corollary 2.3 applied to C_i, and set $C_{i+1} := f_i(T_i) \setminus T_i \subset C_i \setminus T_i$. We stop when we can no longer apply Corollary 2.3, that is, when $|C_{m+1}| < (k-1+\varepsilon)\binom{n}{n/2}$.

We define our fingerprint $(\mathcal{S}, T, C)$ of $\mathcal{G}_k$ by setting

$$T(I) := (T_1, \dots, T_m) \qquad \text{and} \qquad C(T(I)) := C_{m+1} \cup T_1 \cup \cdots \cup T_m,$$

and letting $\mathcal{S} := \{T(I) : I \in \mathcal{I}(\mathcal{G}_k)\}$. Note that Property (iii) implies C is well-defined, while Property (i) guarantees that it is a container function.

In order to check that the constructed fingerprint satisfies the conditions of the theorem, we first bound the sizes of the fingerprints and the number of iterations of the above procedure. To do so, let $2 \leqslant m_0 \leqslant m$ be minimal such that $|C_{m_0}| \leqslant 3k\binom{n}{n/2}$, and observe that, by Property (iv) and the definition (3) of $\tau(A)$,

$$\tau(C_i)|C_i| \leqslant \begin{cases} n^{-3} \cdot 2^n & \text{if } i < m_0, \\ n^{-1} \cdot (1-\delta)^{i-m_0} \cdot 3k\binom{n}{n/2} & \text{otherwise.} \end{cases} \tag{7}$$

The geometric decay of $|C_i|$ moreover immediately implies that $m = O(\log n)$. We thus obtain

$$\sum_{i=1}^{m_0-1} \tau(C_i)|C_i| \leqslant \frac{m \cdot 2^n}{n^3} \ll \frac{1}{n^2}\binom{n}{n/2} \qquad \text{and} \qquad \sum_{i=m_0}^{m} \tau(C_i)|C_i| = \frac{O(1)}{n}\binom{n}{n/2}. \tag{8}$$

Since $|T(I)| = \sum_{i=1}^m |T_i| \leqslant \sum_{i=1}^m k\tau(C_i)|C_i|$, adding the two bounds immediately proves (a). Also, since n is sufficiently large,

$$|C(T(I))| = |C_{m+1}| + |T_1 \cup \cdots \cup T_m| \leqslant (k-1+2\varepsilon)\binom{n}{n/2},$$

which proves (c), since $\varepsilon > 0$ was arbitrary.

It only remains to prove (b), which follows using Lemma 4.3. The first step is to partition the collection of s-sets in $\mathcal{S}$ into subfamilies $\mathcal{S}(\hat{m}_0, \mathbf{t})$, where for given $\hat{m}_0 \in \mathbb{N}$ and $\mathbf{t} = (t_1, \dots, t_{\hat{m}}) \in \mathbb{N}^{\hat{m}}$, we define $\mathcal{S}(\hat{m}_0, \mathbf{t})$ to be set of all $(T_1, \dots, T_{\hat{m}}) \in \mathcal{S}$ such that $\hat{m}_0$ is the smallest integer for which $|C_{\hat{m}_0}| \leqslant 3k\binom{n}{n/2}$ and moreover $|T_i| = t_i$ for each $i \in [\hat{m}]$.

In order to bound the number of elements of $\mathcal{S}(\hat{m}_0, \mathbf{t})$ of size s, set $s_1 = \sum_{i=\hat{m}_0}^{\hat{m}} t_i$, and observe that

$$\sum_{i=\hat{m}_0}^{\hat{m}} t_i \log \frac{1}{t_i} \leqslant s_1 \log \frac{1}{s_1} + \frac{O(1)}{n}\binom{n}{n/2}, \tag{9}$$

by Lemma 4.3 and the second bound in (7). Since each T_i is a subset of the corresponding C_i, we can use the trivial bound $|C_i| \leqslant 2^n$ and the definition of $\hat{m}_0$ to write

$$\begin{aligned}
|\mathcal{S}(\hat{m}_0, \mathbf{t})| &\leqslant \prod_{i=1}^{\hat{m}_0-1} \binom{2^n}{t_i} \prod_{i=\hat{m}_0}^{\hat{m}} \binom{3k\binom{n}{n/2}}{t_i} \\
&\leqslant \left(\prod_{i=1}^{\hat{m}_0-1} 2^{t_i n} \right) \left(\left[3ek \cdot \binom{n}{n/2} \right]^{s_1} \prod_{i=\hat{m}_0}^{\hat{m}} \left(\frac{1}{t_i} \right)^{t_i} \right) \\
&\leqslant \left(\frac{K\binom{n}{n/2}}{s_1} \right)^{s_1} \exp\left(\frac{K}{n} \binom{n}{n/2} \right)
\end{aligned}$$

where the final step follows from the first sum in (8) and from applying the exponential function to (9). Finally, note that the right-hand side is monotone in s_1 on the interval $\big(0, K\binom{n}{n/2}/e\big)$, and we can therefore replace s_1 by s. Summing over the (at most $n^{O(n)}$) choices of $\mathbf{t}$, $\hat{m}_0$ and $\hat{m}$, the claimed bound follows. □

References

[1] J. Balogh, R. Morris, and W. Samotij, Independent sets in hypergraphs, *J. Amer. Math. Soc.*, **28** (2015), 669–709.

[2] J. Balogh, R. Mycroft and A. Treglown, A random version of Sperner's theorem, *J. Combin. Theory, Ser. A*, **128** (2014), 104–110.

[3] B. Bollobás, On generalized graphs, *Acta Math. Acad. Sci. Hungar*, **16** (1965), 447–452.

[4] D. Conlon and W. T. Gowers, Combinatorial theorems in sparse random sets, submitted.

[5] S. Das, W. Gan, and B. Sudakov, Sperner's Theorem and a Problem of Erdős–Katona–Kleitman, *Combin. Prob. Computing*, **24** (2015), 585–608.

[6] P. Erdős, On a Lemma of Littlewood and Offord, *Bull. Amer. Math. Soc.*, **51** (1945), 898–902.

[7] D. Kleitman, A conjecture of Erdős-Katona on commensurable pairs among subsets of an n-set, In *Theory of Graphs (Proc. Colloq., Tihany, 1966)*, pages 215–218. Academic Press, New York, 1968.

[8] Y. Kohayakawa and B. Kreuter, The Width of Random Subsets of Boolean Lattices, *J. Combin. Theory, Ser. A*, **100** (2002), 376–386.

[9] Y. Kohayakawa, B. Kreuter and D. Osthus, The length of random subsets of Boolean lattices, *Random Struct. Algorithms*, **16** (2000), 177–194.

[10] B. Kreuter, Small sublattices in random subsets of Boolean lattices, *Random Struct. Algorithms*, **13** (1998), 383–407.

[11] D. Lubell, A short proof of Sperner's lemma, *J. Combin. Theory*, **1** (1966), 299.

[12] L.D. Mešalkin, A generalization of Sperner's theorem on the number of subsets of a finite set, *Teor. Verojatnost. i Primenen*, **8** (1963), 219–220.

[13] R. Morris and D. Saxton, The number of $C_{2\ell}$-free graphs, to appear in *Adv. Math.*

[14] D. Osthus, Maximum antichains in random subsets of a finite set, *J. Combin. Theory, Ser. A*, **90** (2000), 336–346.

[15] A. Rényi, On random subsets of a finite set, *Mathematica (Cluj)*, 355–362, 1961.

[16] D. Saxton and A. Thomason, Hypergraph containers, *Inventiones Math.*, **201** (2015), 925–992.

[17] M. Schacht, Extremal results for random discrete structures, submitted.

[18] E. Sperner, Ein Satz über Untermengen einer endlichen Menge, *Math. Z.*, **27** (1928), 544–548.

[19] K. Yamamoto, Logarithmic order of free distributive lattice, *J. Math. Soc. Japan*, **6** (1954), 343–353.

IMPA, Estrada Dona Castorina 110, Jardim Botânico, Rio de Janeiro, RJ, Brasil
E-mail address: {collares|rob}@impa.br